

Calibers: A Bandwidth Calendaring Paradigm For Science Workflows

Fatemah Alali^{1,3}, Nate Hanford^{2,3}, Eric Pouyoul³, Raj Kettimuthu⁴, Mariam Kiran³, Ben Mackcrane⁵, Brian Tierney³, Yatish Kumar⁵ and Dipak Ghosal^{2,3}

¹Department of Electrical and Computer Engineering, University of Virginia, Virginia, USA. Email: fha6np@virginia.edu

²Department of Computer Science, University of California, Davis CA, USA. Email: nate.hanford@gmail.com, dghosal@ucdavis.edu

³Energy Sciences Network, Lawrence Berkeley National Laboratory, California, USA, Email: lomax@es.net, mkiran@es.net, bltierney@es.net

⁴Mathematics and Computer Science Division, Argonne National Laboratory, Chicago, USA. Email: kettimut@anl.gov

⁵Corsa Technology. Email: ben.mackcrane@corsa.com, yatish.kumar@corsa.com

Abstract

Many scientific workflows require large data transfers between distributed instrument facilities, storage and computing resources. To ensure that these resources are maximally utilized, R&E networks connecting these resources, must ensure that there is no bottleneck. However, running the network at high utilization often results in congestion and poor end-to-end TCP throughput performance and/or fairness. This in turn leads to unpredictability in transfer time and poor utilization of distributed resources. Calibers (Calender and Large-scale Bandwidth Event-driven Simulations) aims to advance state-of-the-art in traffic engineering by leveraging SDN-based network architecture and flow pacing algorithms to provide predictable data transfers performance and higher network utilization. Calibers highlights how by intelligently and dynamically shaping flows, we can maximize the number of flows that achieve deadline while improving network resource utilization.

In this paper, we present a prototype architecture for Calibers that uses a central controller with distributed agents to dynamically pace flows at the ingress of the network to meet deadlines. Using Globus/Grid-FTP, we experimentally demonstrate that pacing can be used to meet data transfer deadlines which cannot be achieved using TCP. Finally, we present dynamic flow pacing algorithms that maximize acceptance ratio of flows for which deadlines can be met while maximizing network utilization. Our results show that simple heuristics that optimizes locally on the most bottlenecked link can perform almost as well as heuristics that attempt to optimize globally.

Keywords: Bandwidth Calendaring, Flow Pacing, Software Defined Network (SDN), TCP, Dynamic Flow Pacing, Simulation Analysis, Traffic Shaping

1. Introduction

Scientific analysis in experiments such as high-energy physics or climate modeling, usually involve extremely complex workflows to ensure successful and reliable results. These workflows include a number of tasks, involve multiple actors, software and infrastructures, that work together as a workflow from data generation to delivery. For example, in the Advanced Light source (ALS) data is generated from multiple detectors which is then collected on an NERSC supercomputing data center via high-speed network connections. It is imperative that the data is delivered in a timely manner, with minimum loss, such that further computations can be performed using supercomputing resources that have to be *a priori* reserved. In order that the supercomputing resources are maximally utilized, this requires the net-

work service to allow deadlines for large data transfers.

There are two approaches to ensure that the data transfers can be made with predictable performance and within requested deadlines. One approach is to use advanced reservations of links, such as OSCARS or open NSA [1], that allow setting up circuits of specified capacities between routers. Advanced reservation schemes require additional time to setup circuits, are only associated with WAN border routers and are difficult to automate due to required user knowledge, network topology and request details. Furthermore, applications do not generate traffic all the time which leads to wasted reserved capacity.

The second approach is to run the network at low utilization and use standard TCP. New TCP protocols, such as TCP Hamilton [2] and BBR-TCP [3], can effi-

ciently adapt to the bottleneck capacity and where multiple competing flows are involved, they *equally* split the bottleneck capacity. However, even with the new TCP algorithms, sustained bottlenecks lead to unpredictable throughput performance and difficulties in arbitrarily splitting bottlenecked bandwidths among competing flows. Finally, as the growth in data transfer volume out-paces the increase in the data link rates, running the network at low utilization is not cost effective [4].

To help accelerate the effort to run the network at high utilization and enable deadline aware data transfers, network automation through Software-Defined Networks (SDN) are being advanced to control network traffic depending on data demand. In principle, SDN allow individual switches to be managed and controlled following centralized traffic engineering principles [5]. Furthermore, SDN switches provide the ability to pace traffic at ingress of the network. These features in addition to TCP protocol, or the pacing algorithm at the source nodes [3], together provide the necessary tools to dynamically allocate bandwidth to flows for meeting deadlines while ensuring the network operates at high utilization [6, 7].

This paper aims to implement a centralized traffic engineering approach and control distributed agents at the edge (ingress point of the network) to dynamically pace flow for meeting transfer deadlines, while achieving high network utilization. The dynamic pacing algorithm is able to analyze traffic patterns and follow a rolling horizon model to pace flows at appropriate rates to optimize network performance and meet deadlines. As a result Calibers, not only calendars flow, but also lays the foundation for future work where these capabilities can be coupled with advanced tools to control networks dynamically.

Following are the main contributions of this work:

1. We describe an architecture that implements bandwidth calendaring for scientific workflows. The architectures leverage SDN switches that can pace flows at the ingress point. The architecture implements a central controller with distributed agents at the edge of the network that monitor flow performance and implement dynamic flow pacing set by the controller.
2. We present experimental results using Globus and GridFTP that show the importance of pacing in achieving deadline aware data transfer service. We compare our results with TCP Hamilton results.
3. We propose different heuristic algorithms based on combining two orthogonal principles - 1) local vs

global optimization and 2) Shortest Job First vs Longest Job First (LJF). We perform a preliminary performance comparison of these algorithms with respect to a performance metric efficacy that is defined as the ratio of the percentage of request accepted to the network utilization. Our results show that simple heuristics, that optimize locally on the most bottlenecked link can perform almost as well as heuristics that attempt to optimize globally.

The remainder of this paper is organized as follows. In Section 2, we discuss the motivation of our work specifically the the importance of deadline aware data transfers in scientific workflows. We also discuss the importance of pacing and traffic shaping in deadline-aware traffic flow-scheduling. In Section 3, we present the architecture of Calibers in a software defined network. In Section 4 we present experimental results on a preliminary prototype Calibers architecture to demonstrate the effectiveness in meeting deadlines compared standard TCP. In Section 5 we present work on a dynamic flow pacing algorithm and present preliminary simulation results. In Section 7, we present the related work followed by conclusions and future work in Section 8.

2. Motivation

It is often the case that a large data transfer is inherently deadline-aware. For example, an HPC user may want to ensure that data is present at an HPC site before conducting their experiments. However, with the unpredictability of network resource utilization, this can pose a problem. Latency variation coupled with TCP's typical "sawtooth" behavior can lead to a lack of predictability in meeting deadlines. Furthermore, even when TCP achieves pareto-optimality, this behavior may not always be desired. For example, if one flow must complete faster over a bottleneck link than others, simply fairly sharing the available bottleneck bandwidth may not achieve the desired deadline goal.

Neuman et al. [8] highlighted the need to redesign I/O architecture and network links to cope with the performance of distributed science instruments. Salsano et al. [9] discussed various APIs to optimize packet movement based on user information to improve quality of experience. Also, OpenStack clouds have been investigating how workload can be balanced over geographically distributed data centers. However, to the best of our knowledge, there is still a lack of an implementation system which can demonstrate dynamic traffic shaping [4]. This paper aims to contribute to this research area

by studying how flows can be dynamically paced to reduce loss and optimize link performance.

Towards, the above goal, we first demonstrate that even in a controlled, isolated environment, TCP does not provide predictable data transfer rate in the event of congestion. By pacing each data transfer properly, network congestion can be avoided and flow completion can be predicted. The following sections discuss the strategies and algorithm utilized to decide rate flows for transmitting data and how to achieve maximum number of successful flows with network utilization.

2.1. Pacing and Shaping Traffic

Active Queue Management (AQM) has existed on core network routers and switches for decades. In particular, Fair Queuing with Controlled Delay (FQ CoDel) has been implemented in such switches and routers for decades and was introduced in the Linux kernel in 2012 [10]. FQ implementations typically work on the principle of creating some data structure of ongoing flows, and then using a Deficit Round Robin (DRR) approach in order to dequeue packets from their respective buffers. In this fashion, a lightweight, but reasonably "fair" allocation of bandwidth resources occurs between multiple competing flows. CoDel implementations essentially work on the principle of putting hard limits on the real queue size of ongoing flows such that a particular latency target is met. In such a fashion, the *delay* of a flow is *controlled*, which has important implications for minimizing standing queue sizes elsewhere in the network and therefore delivering more predictable performance [11]. However, CoDel can also be used in conjunction with FQ to pace flows under a particular maximum rate without violating fairness principles. This is precisely how we conduct end-system pacing in Calibers.

3. Architecture

Workflow orchestrators provision resources across the network, with assumptions that it does not introduce performance penalties. Calibers is an experimental network service, targeted to higher level resource orchestrators. It focuses on optimizing network resources such that each data flows (i.e. file transfers) performs at least at the minimum average rate over the transfer duration. This allows Calibers to provide deadline delivery guarantees.

The experimental software platform designed in Calibers involves several components:

- A REST/JSON API: Orchestrators use this API to schedule file transfers. They provide source, destination, file size, deadline and maximum I/O rate of the endpoints of the transfer.
- An event publisher: Allows orchestrators to obtain real-time information on the maximum rate the network has allocated to data transfers without experiencing network congestion.
- SDN-based rate shaper: This component enforces flows to not exceed their bandwidth allocation.
- Calibers optimization algorithm: This dynamically adjusts the maximum rate of each flow, ensuring that all flows are on track to meet their respective deadlines. This in turn increases network utilization and maximizes the request admission rate.

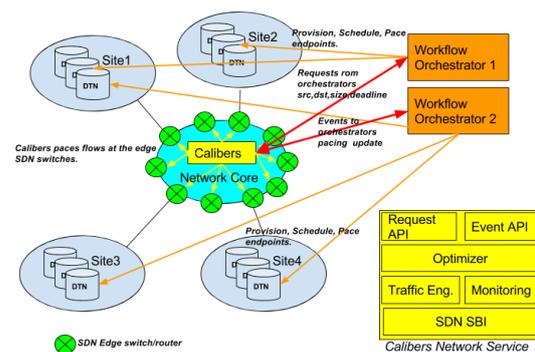


Figure 1: Calibers architecture illustrating the various components.

Calibers aims to experiment with higher level of network services, file transfer deadlines and demonstrate new paradigms to higher network utilization. It makes several assumptions to realize these goals. These are not always true in a production environment. For example, Calibers assumes that endpoints are sufficiently provisioned with I/O, networking and processing resources. Another important assumption is that Calibers is given a minimum guaranteed capacity on the overall network and therefore prevents it to be impacted by non Calibers traffic.

When the flow pacing rate is dynamically changed at the network edge, it may result in packet loss which may result in throughput loss. Note that TCP algorithms such as H-TCP [2] and BBR-TCP [3] can quickly adapt to these changes. Additionally, new source pacing algorithms based on model predictive control [12] can also

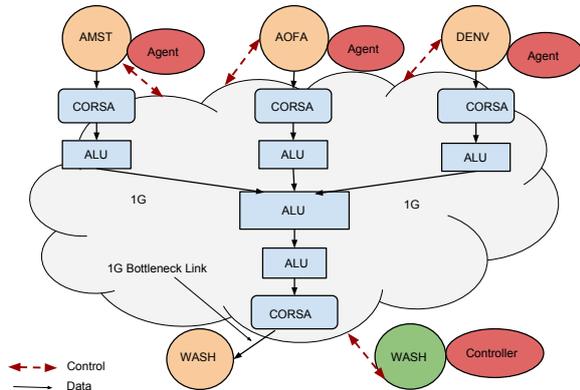


Figure 2: The experimental setup. Data transfer are initiated from Amsterdam (AMST), New York (AOFY) and Denver (DENV) to Washington (WASH). The Controller in a different VM in WASH can interact with the distributed agents to pace the flows.

be used to ensure that the source quickly adapts to edge pacing rates.

4. Experimental Results

The experimental setup in Figure 2 is based on ESnet’s SDN Testbed, a high-speed Wide-Area Network (WAN) SDN-ready testbed spanning two continents. The backbone data rates are guaranteed and setup via dedicated OSCARS circuits.

The testbed closely resembles ESnet’s high-speed production network in both hardware and topology, as it is an overlay of the ESnet production WAN [13]. Using ProxMox, a Linux container management system, we define three senders Amsterdam (AMST), New York (AOFY) and Denver (DENV) with varied round-trip latencies to one receiver at Washington (WASH). A controller container was provisioned at Washington (WASH). In order to simulate a real workflow, we use real FTP file transfers with 10 gigabyte data files rather than benchmarking tools. In particular, we used the Globus GridFTP file transfer tool, with the Globus API running on the controller, to synchronize the start of all three transfers [14]. Furthermore our topology closely simulates the case of a bottlenecked receiver obtaining data from three different senders with different geographical locations.

Our preliminary research was focused on determining the feasibility of pacing both at the network’s edge with SDN-enabled Corsas switches [15] and within the end-system itself. In this fashion, the Corsas switches act to police bandwidth utilization below a certain threshold.

Although the method of restricting bandwidth utilization is indeed fair shaping, this is important for preservation the overall fairness. However, in our current implementation, losses at the Corsas switch, due to limits imposed on the queue size, will result in activation of the H-TCP congestion avoidance algorithm. In order to avoid this, pacing at the end-system is used. One could imagine in an active deployment that an end-user would either participate in pacing, or at least be policed and therefore unable to interfere with the pacing of other senders to a bottleneck link. Thus, the orchestrator is capable of limiting interference and preserving the overall workload allocation of the system.

Before we tested the feasibility of pacing with Globus and GridFTP, we conducted experiments testing the effects of pacing on nuttcp flows sharing a bottleneck link. In Figure 3, the throughput performance of one of those bottlenecked flows is shown both in the paced and unpaced scenarios.

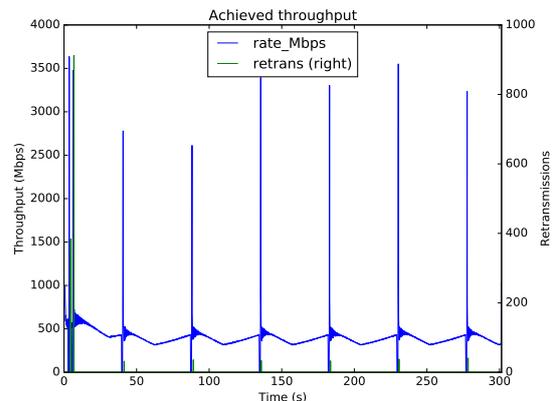


Figure 3: This experiment considers two H-TCP flows sharing a 1 Gbps link with both flows are un-paced. The spikes correspond to the buffer build up prior to loss events.

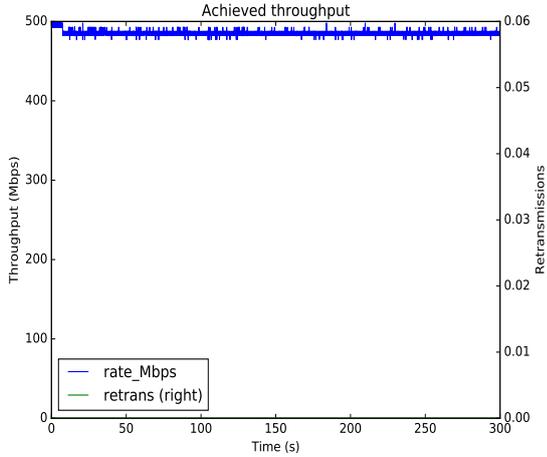


Figure 4: This experiment considers two H-TCP flows sharing a 1 Gbps link and both flows are paced at 500 Mbps.

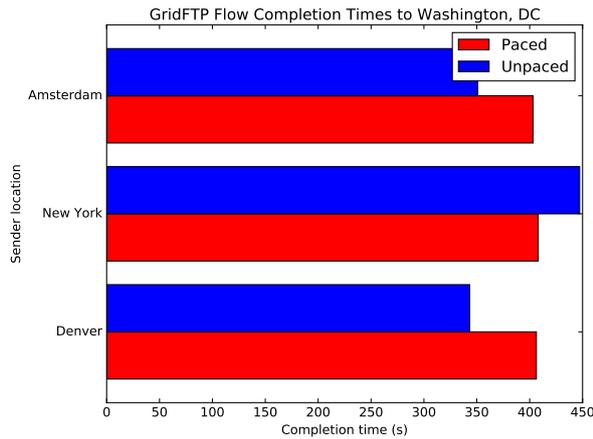


Figure 5: Flow completion times of three hosts sharing a bottleneck link.

In Figure 5, the difference in flow completion times for 3 equal-size, equally-paced files can be compared with the unpaced flow completion times for the same equal-size transfers. Here, the feasibility of setting reasonable deadlines with the help of end-system pacing is demonstrated. The unpaced flows do not end up sharing the bottleneck equally for their duration, and the inherent unpredictability of unpaced TCP at high throughputs given different latencies is apparent. The paced flows, however, all complete within two seconds of each other, demonstrating that pacing can work effectively with the Globus GridFTP architecture.

5. Scheduling Algorithm for Dynamic Pacing

As mentioned before, the objective of the scheduler is to decrease the number of rejected data transfer requests while increasing the network utilization. The notations used is shown in Table 1. The scheduler operates at fixed discrete epochs with the following assumptions: (i) each link has a free capacity I_i^C to be used by the scheduler, (ii) start time for each data transfer request is immediate, i.e., the scheduler does not support advance reservation, and (ii) the scheduler updates the network status periodically every scheduling interval (epoch).

The scheduling problem is divided into to sub-problems: (i) *new flow*: when a new request arrives, how to decide whether to accept or reject the request? and (ii) *completed flow*: when a request completes, how to distribute the free capacity among the ongoing flows? We study four heuristic algorithms by combining two concepts: (i) global and local optimization and (ii) Shortest Job First (SJF) and Longest Job First (LJF).

In the global approach, the scheduler consider all the flows when distributing any residual capacity. On the other hand, the local approach focus on the bottleneck links in the network and distribute the residual capacity by reallocating locally only the flows that span the bottleneck link(s). The LJF and SJF are known concepts where longest jobs are favored with LJF and shorts jobs are favored when SJF is used. This concept is used by both the global and local scheduler in the following way. When the scheduler decides (locally or globally) which flows should be considered when distributing the residual capacity, SJF or LJF will be used to decide the order in which the flows will be assigned the residual capacity.

The scheduler keeps track of multiple parameters as shown in Table 1. One of the most important parameter the scheduler uses to make decision is $R_{f_i}^{min}$, which is the minimum required rate to ensure the flow will not miss its deadline. The pseudo-code of both the global and local-approach is provided in Alg. 1. Both schedulers use the same approach when a new request arrives, however, they differ in the way the capacity is redistributed when a request completes.

5.1. Approach 1: Global Optimization

Sub-problem 1: *new flow*, when a new flow f_i corresponding to request u_i arrives, the scheduler computes $R_{f_i}^{min}$, and checks if $R_{p_i}^{resid}$ is greater than $R_{f_i}^{min}$, it assign $R_{p_i}^{resid}$ to the new flow. The scheduler gives the maximum available rate to the new request instead of giving it $R_{f_i}^{min}$ for two reasons: (i) to increase the link utilization, and (ii) to complete the file transfer as soon as possible in order to free up the resources to accept future

requests. If $R_{f_i}^{min}$ is not available, the scheduler move to the second phase, which is pacing other flows in order to accept the new flow.

In the pacing phase, for each link l_i , the scheduler finds the list of flows I_i^F span link l_i . SJF or LJF concepts are used to decide which flow(s) of the list I_i^F to pace (or slow down). When using SJF, the scheduler favors short flows with longest flows being slowed down first and vice versa. The scheduler paces the first flow in the sorted list by taking its slack rate $R_{f_j}^{slack}$ and assigns it to the new flow. If the first flow slack ($R_{f_j}^{slack}$) is less than the new flow required minimum rate ($R_{f_i}^{min}$), then the scheduler takes the slack rate of the second flow in the sorted list until the sum of the slack rates is equal to the new flow required minimum rate, or until there are no more flows in the sorted list. Hence, the request will be rejected because even with pacing $R_{f_i}^{min}$ cannot be assigned to the new flow.

Sub-problem 2: *completed flow*, at the beginning of each epoch, the scheduler checks if a flow has completed by checking the flow completion time t_c . The flow completion time is a dynamic parameter, that changes based on the allocated rate ($R_{f_i}^{alloc}$). For all the flows completed at the scheduling epoch, the scheduler traverse the path of each completed flow and finds the set of other flows, that span the links in the path (involved_flows). After finding all involved flows, the scheduler now has a global view and starts distributing the residual capacity using SJF or LJF concepts.

5.2. Approach 2: Local Optimization

As mentioned earlier, the same approach is used when a new request arrives. However, the local scheduler takes a different approach when a flow completes.

Sub-problem 2: *completed flow*, at the begging of each epoch, the scheduler checks if a flow has completed. For each completed flow (f_i), the scheduler finds the bottleneck link ($l_{bottleneck}$) in p_i . The link which has a flow with the maximum t_c , is the link that will stay busy the longest, hence, is the bottleneck link that might cause future requests to be rejected. Therefore, by freeing up only this link the probability of accepting flows in the future increases. The scheduler considers only the flows spanning the bottleneck link when distributing the residual capacity, which in contrast to the global approach, where scheduler considers all flows spanning all links of all completed flows paths.

6. Simulation Analysis

Flow-level simulation was conducted to evaluate the performance of the four schedulers: (i) local-SJF, (ii)

local-LJF, (iii) global-SJF, and (iv) global-LJF. The simulator was written using Python and each simulation was executed until 30k requests were generated.

6.1. Simulation Setup

Network:. Google's inter-data center network G-scale [5] with 12 nodes and 19 links was used to evaluate. The links capacity l_i^C was set to 10 Gbps for all links in the network.

Workload:. Requests were generated as follow: (i) Request inter-arrival time was modeled with an exponential distribution with arrival rate λ varying between 0.05 to 1.6 with step of 0.1, i.e. the mean inter-arrival time between requests varies from 20 sec to 0.625 sec. (ii) Request deadline time was modeled following an exponential distribution with average deadline (t_d) of 1 hour. (iii) As the file size is related to the deadline, it was modeled as follows [7]. First, a transfer rate (i.e. $R_{f_i}^{min}$) is sampled following an exponential distribution with average rate of 100 Mbps. Next the file size was computed as transfer rate $\times t_d$. This results in a product distribution with a mean file size of 45 GB. (iv) Source and destination pairs were picked uniformly.

Metrics. : Three performance metrics were measured: (i) Network utilization is computed by measuring the link utilization per second for each link in the network $L_i^{utilization}(t)$, then taking the average utilization for each link across the entire simulation time ($\forall l_i \in L, L_i^{utilization} = \text{mean}(L_i^{utilization}(t))$). Finally, the network utilization is measured as the average of $L_i^{utilization}$ for all the links in the network (network utilization = $\text{mean}[L_i^{utilization} \forall l_i \in L]$). (ii) Reject ratio which is defined as the number of rejected requests divided by the total number of requests. (iii) Performance index is defined as the difference between network utilization and reject ratio. The larger the difference, the better is the performance of the scheduler as we want an ideal 100% utilization and a reject ratio of 0%.

6.2. Results

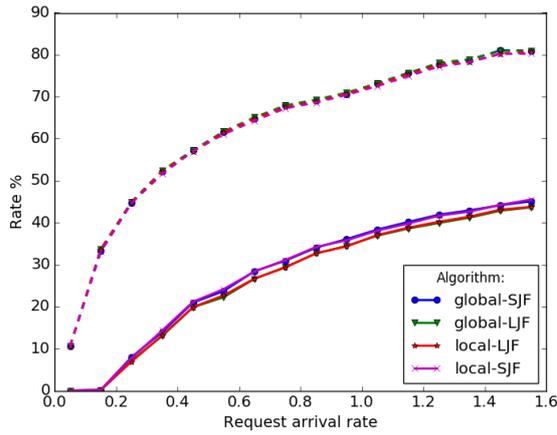
Figure 6 shows the scheduler performance for the G-scale network with mean file size of 45 GB (transfer rate of 100 Mbps). The performance difference between the local and global approach is negligible. This shows that redistributing the capacity only in the bottleneck link of the path of the completed flow, is enough to perform good, as when considering all flows traversing the path. Also, Fig. 6 shows that LJF is performing slightly better in both local and global approaches. LJF reduces the

Algorithm 1: Dynamic Pacing Scheduler

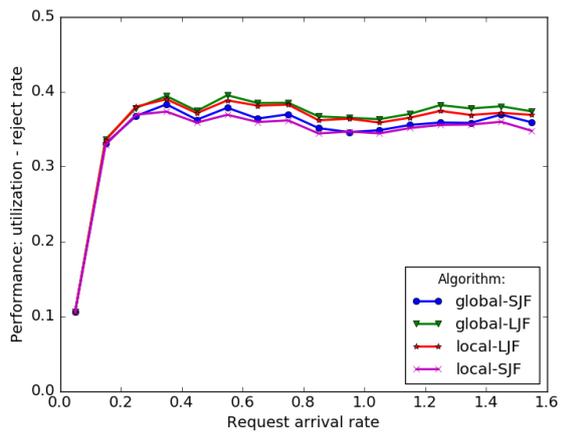
```

Input:  $U$ 
1  remove_completed_flows( $t_{now}$ )
2  foreach  $u_i \in U$  do
3     $R_{p_i}^{resid} = \min(R_{l_i}^{resid} \forall l_i \in p_i)$ 
4    if  $R_{p_i}^{resid} < R_{f_i}^{min}$  then
5      pace( $f_i$ )
6    else
7       $R_{f_i}^{alloc} = \min\{R_{f_i}^{max}, R_{p_i}^{resid}\}$ 
8      return success
9  Function pace( $f_i$ )
10   $\forall l_i \in p_i$ 
11   $R_{temp} = 0$ 
12   $R_{temp}^{p_i} = []$ 
13  found = False
14  involved_flows =  $l_i^F$ 
15  sorted_involved_flows = sort the list in ascending (if
16  LJF) or descending (if SJF) order based on  $t_c$ 
17  foreach  $f_j \in sorted\_involved\_flows$  do
18     $R_{temp} = R_{temp} + R_{f_j}^{slack}$ 
19    if  $R_{temp} \geq R_{f_i}^{min}$  then
20      found = True
21      add  $R_{temp}$  to  $R_{temp}^{p_i}$ 
22      break
23  if found = True then
24     $R_{f_i}^{alloc} = \min(R_{temp}^{p_i})$ 
25    return success
26  else
27    return reject
27  Function remove_completed_flows()
28  involved_flows = []
29  involved_links = []
30  foreach  $f_i \in F$  do
31    if  $t_c = t_{now}$  then
32       $\forall l_j \in p_i$ 
33      remove  $f_i$  from  $l_j^F$  and  $F$ 
34      add  $l_j$  to involved_links
35      add to  $l_j^F$  involved_flows
36    if local-sched then
37      local_reshape(involved_links)
38  if global-sched then
39    global_reshape(involved_flows)
40  Function global_reshape(involved_flows)
41   $R_{f_i}^{alloc} = R_{f_i}^{min} \forall f_i \in involved\_flows$ 
42  sorted_involved_flows = sort the list in ascending (if
43  SJF) or descending (if LJF) order based on  $t_c$ 
44  foreach  $f_i \in sorted\_involved\_flows$  do
45     $R_{p_i}^{resid} = \min(R_{l_i}^{resid} \forall l_i \in p_i)$ 
46     $R_{f_i}^{alloc} = R_{f_i}^{alloc} + R_{p_i}^{resid}$ 
46  Function local_reshape(involved_links)
47  find  $l_{bottleneck}$ 
48   $R_{f_i}^{alloc} = R_{f_i}^{min} \forall f_i \in l_{bottleneck}$ 
49  involved_flows =  $l_{bottleneck}^F$ 
50  sorted_involved_flows = sort the list in ascending (if
51  SJF) or descending (if LJF) order based on  $t_c$ 
52  foreach  $f_i \in sorted\_involved\_flows$  do
53     $R_{p_i}^{resid} = \min(R_{l_i}^{resid} \forall l_i \in p_i)$ 
54     $R_{f_i}^{alloc} = R_{f_i}^{alloc} + R_{p_i}^{resid}$ 

```



(a) Reject ratio and network utilization. Solid lines represent reject ratio and dashed lines represent network utilization



(b) Performance index

Figure 6: Performance comparison of the four algorithms for the G-scale network.

Table 1: Notation used.

Symbol	Description
t_{now}	Current scheduling epoch
U	set of requests at time
u_i	User i request which is defined by 5 tuples $(IP_{src}, IP_{dst}, S, t_d)$ where IP_{src} and IP_{dst} are the source and destination IP addresses, S is the size of data in Mbytes, and t_d is the deadline
F	The set of the currently scheduled flows
f_i	flow $i \in F$
p_i	A path p_i is sequence of links corresponds to f_i source and destination pair
$R_{f_i}^{min}$	For flow f_i , this is the minimum rate that will guarantee that the deadline is met. This is $S/(t_d - t_s)$
$R_{f_i}^{alloc}$	For flow f_i , the pacing rate that was allocated for scheduling epoch
$R_{f_i}^{slack}$	The slack rate assigned to flow f_i , $R_{f_i}^{slack} = R_{f_i}^{alloc} - R_{f_i}^{min}$
L	The set of the links in the network
l_i	Link i in the network, $l_i \in L$
l_i^C	Link l_i capacity
l_i^F	List of flows span link l_i
$R_{l_i}^{resid}$	The residual capacity for link l_i . $R_{l_i}^{resid} = l_i^C - \sum_{f_i \in l_i^F} R_{f_i}^{alloc}$
$l_{bottleneck}$	The link with the flow that has the maximum t_c

makespan time of all flows by assigning more rate to the flows with the longest completion time. This results on freeing up the links faster to accommodate future requests. On the other hand, SJF frees up some capacity of the link earlier than LJF but the makespan of the flows stays the same.

The difference in performance between SJF and LJF becomes more apparent with a longer epoch duration as shown in Fig. 7. Many observations can be made with the longer epoch duration. First, the network utilization for LJF is higher compared to the SJF with the epoch of 5 mins. Since SJF favors the flows with the lowest completion time when redistributing the residual capacity, then the probability of the flow finishing during the epoch is higher compared to if LJF is used. Therefore the flow will finish during the epoch and the capacity used by the completed flow will be wasted as no reshaping is done within the epoch. Second, the reject rate increases as the epoch duration increases, as the requests are aggregated making the scheduler less flexible. Third, at a lower arrival rate, the utilization with the epoch duration of 5 mins is higher than the utilization with epoch of 1 sec.

Lets consider the lowest arrival rate which is 0.05, i.e. the inter-arrival between requests is 20 sec. With this rate an average of 15 requests are aggregated and passed to the scheduler at the beginning of the scheduling epoch. The scheduler will assign lower rates for the requests since there are many of them. Hence, the flows

will take longer to finish and the links will stay busy (utilized) longer. On the other hand, with an epoch of 1 sec, whenever a request comes, the scheduler gives it a maximum available rate making the flow completion time small. Hence, the flow completes early and links stay idle.

7. Related Work

The over-arching goal of this work is to deliver deadline aware data transfers as a network service, while ensuring high network utilization. We leveraged SDN with the ability to perform dynamic traffic pacing at the network edge. There are a number of recent studies with similar goals. In the following paragraphs, we review the related work and point out the key differences from our work.

There has been a number of prior studies on flow pacing [16, 17, 18]. Broadly speaking, flow pacing can be performed at the source host or at the edge where the access network connects to the core network. The former is referred to as host pacing, or more commonly TCP pacing, while the later is referred to as edge pacing and can be performed by the network service provider [18]. In this paper, we study edge pacing enabled softwerized SDN switches.

SDN networks allow dynamic and centralized traffic engineering (TE) via flow pacing. B4 [5] presents Google's effort in leveraging SDN to centralized TE

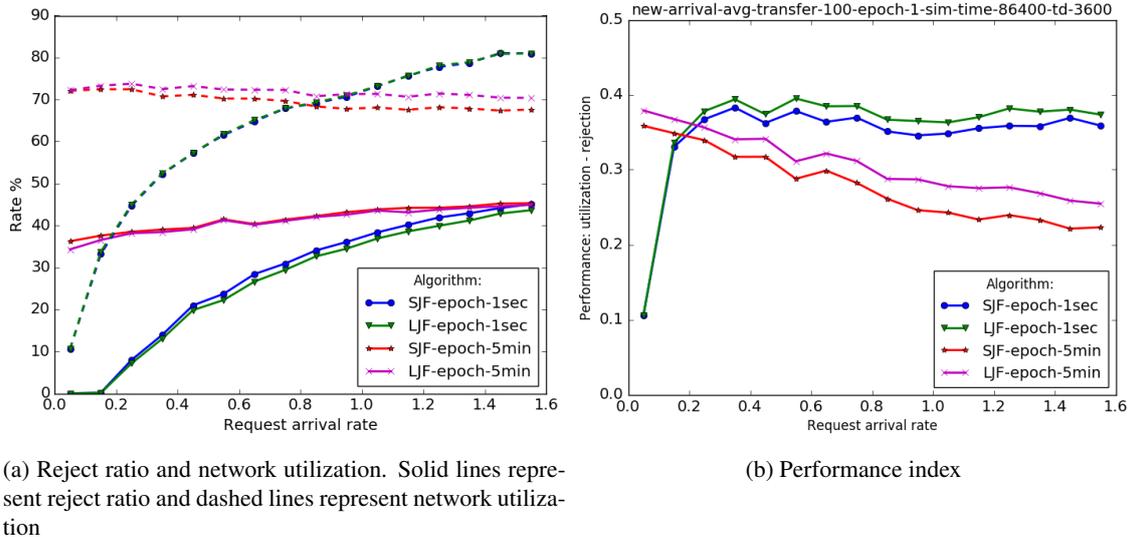


Figure 7: The impact of epoch duration on the performance.

and drive links to near full utilization. As similar study SWAN [19] also improves network utilization of inter-DC WAN by scheduling the service traffic in a centralized manner. However, all of these studies do not consider deadline associated with their transfers. The study in Tempus [6] considered deadlines and developed an optimization framework to maximize the fraction of transfer delivered before deadline, ensuring fairness among all requests. This work, however also does not guarantee meeting deadlines.

In a recent study [7], deadlines have been investigated in the context of inter-data center data transfers. Building on a deadline aware network abstraction (DNA) where transfer deadlines can be specified, the study proposes AMEOPA which uses traffic shaping at the source, to meet data transfer deadlines. While this study is the most similar to ours, there are a few key differences. First, this study focuses on scientific workflows. Second, we consider edge pacing. Finally, we show that simple dynamic pacing algorithms that optimizes locally on the most bottleneck link perform as well more complex algorithms that attempt to optimize globally.

8. Conclusions and Future Work

Calibers has demonstrated, in an ideal situation of a controlled environment, that TCP congestion avoidance algorithms, while performing well at maximizing network utilization, it cannot provide the desired behavior for workflow orchestration. In particular, TCP relies

on network characteristic, such as RTT, packet retransmission, pace flows and ignore flow needs. As result some flows may go faster than they need, while others may go slower than they should, such to meet deadlines and maximize resource utilization. However, the performance of modern TCP, in conjunction with the ability of SDN to implement centralized traffic engineering, allows Calibers to optimize network utilization to provide predictable performance. Our preliminary study on dynamic pacing algorithms suggests that simple heuristics that optimize locally on the most bottlenecked link can perform almost as well as attempts to optimize globally.

In our future work, we will address some of the assumptions that we made in this preliminary work. In particular, we will consider the interaction between deadline aware traffic and background traffic. We will investigate methods to predict the background traffic and leave appropriate network capacity to minimize the impact on the deadline aware traffic. This work will also enhance Calibers to accept deadline aware data transfer request at a future time.

- [1] ESnet, Oscars: On-demand secure circuits and advance reservation system <https://www.es.net/engineering-services/oscars/>.
- [2] D. Leith, R. Shorten, H-tcp: Tcp congestion control for high bandwidth-delay product paths, draft-leith-tcp-htcp-06 (work in progress) (2008).
- [3] N. Cardwell, Y. Cheng, C. S. Gunn, S. H. Yeganeh, V. Jacobson, Bbr: Congestion-based congestion control, Queue 14 (2016) 50.
- [4] DOE, Doe network 2025: Network research problems and challenges for doe scientists workshop (2016).
- [5] S. Jain, A. Kumar, S. Mandal, J. Ong, L. Poutievski, A. Singh, S. Venkata, J. Wanderer, J. Zhou, M. Zhu, J. Zolla, U. Hölzle,

- S. Stuart, A. Vahdat, B4: Experience with a globally-deployed software defined wan, *SIGCOMM Comput. Commun. Rev.* 43 (2013) 3–14.
- [6] S. Kandula, I. Menache, R. Schwartz, S. R. Babbula, Calendar- ing for wide area networks, in: *Proceedings of the 2014 ACM Conference on SIGCOMM, SIGCOMM '14*, ACM, New York, NY, USA, 2014, pp. 515–526.
- [7] H. Zhang, K. Chen, W. Bai, D. Han, C. Tian, H. Wang, H. Guan, M. Zhang, Guaranteeing deadlines for inter-data center trans- fers, *IEEE/ACM Trans. Netw.* 25 (2017) 579–595.
- [8] H. Newman, A. M. I. Legrand, R. Voicu, D. Kcira, J. Bunn, High speed scientific data transfers using software defined net- working, in: *Innovating the Network for Data Intensive Science (INDIS)*, SC 15.
- [9] S. Salsano, N. Blefari-Melazzi, A. Detti, G. Morabito, L. Veltri, Information centric networking over sdn and openflow: Archi- tectural aspects and experiments on the ofelia testbed, in: *Int. Jour. Comp. and Tele. Net. archive*, 57(16), 3207-3221.
- [10] D. Raghuvanshi, B. Annappa, M. Tahiliani, On the effectiveness of codelet for active queue management, in: *Advanced Comput- ing and Communication Technologies (ACCT), 2013 Third Inter- national Conference on*, pp. 107–114.
- [11] J. Gettys, K. Nichols, Bufferbloat: Dark buffers in the internet, *Queue* 9 (2011) 40:40–40:54.
- [12] D. Fridovich-Keil, N. Hanford, M. Chapman, C. Tomlin, M. Farrens, D. Ghosal, A model predictive control approach to flow pacing for tcp, in: *55th Annual Allerton Conference on Communication, Control, and Computing*.
- [13] Energy Sciences Network, ESnet 100G SDN Testbed <http://es.net/network-r-and-d/experimental-network-testbeds/100g-sdn-testbed>.
- [14] W. Allcock, J. Bresnahan, R. Kettimuthu, M. Link, C. Du- mitrescu, I. Raicu, I. Foster, The globus striped gridftp frame- work and server, in: *Proceedings of the 2005 ACM/IEEE Con- ference on Supercomputing, SC '05*, IEEE Computer Society, Washington, DC, USA, 2005, pp. 54–.
- [15] CORSAs, SDN switches, <http://www.corsa.com/solutions>.
- [16] A. Aggarwal, S. Savage, T. Anderson, Understanding the perfor- mance of tcp pacing, in: *INFOCOM 2000. Nineteenth Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings. IEEE*, volume 3, IEEE, pp. 1157–1165.
- [17] D. Wei, P. Cao, S. Low, C. EAS, Tcp pacing revisited, in: *Proceedings of IEEE INFOCOM*.
- [18] H. H. Gharakheili, A. Vishwanath, V. Sivaraman, Comparing edge and host traffic pacing in small buffer networks, *Computer Networks* 77 (2015) 103 – 116.
- [19] C.-Y. Hong, S. Kandula, R. Mahajan, M. Zhang, V. Gill, M. Nanduri, R. Wattenhofer, Achieving high utilization with software-driven wan, in: *Proceedings of the ACM SIGCOMM 2013 Conference on SIGCOMM, SIGCOMM '13*, ACM, New York, NY, USA, 2013, pp. 15–26.